

**ARMSTRONG
LABORATORY**

AL/HR-TR-1995-0194



**IMAGE UPDATE RATE
CAN AFFECT THE PERCEIVED SPEED
OF SIMULATED SELF-MOTION**

**Julie Mapes Lindholm
Timothy M. Askins
Norwood Sisson**

**Hughes Training, Inc., Training Operations
6001 S. Power Road, Bldg 561
Mesa, AZ 85206-0904**

**HUMAN RESOURCES DIRECTORATE
AIRCREW TRAINING RESEARCH DIVISION
6001 S. Power Road, Bldg 558
Mesa, AZ 85206-0904**

June 1996


19961106 145

Final Technical Report for Period June 1992 to May 1995

Approved for public release; distribution is unlimited.

**AIR FORCE MATERIEL COMMAND
BROOKS AIR FORCE BASE, TEXAS**

DISCLAIMER NOTICE



**THIS DOCUMENT IS BEST
QUALITY AVAILABLE. THE COPY
FURNISHED TO DTIC CONTAINED
A SIGNIFICANT NUMBER OF
PAGES WHICH DO NOT
REPRODUCE LEGIBLY.**

NOTICES

When Government drawings, specifications, or other data are used for any purpose other than in connection with a definitely Government-related procurement, the United States Government incurs no responsibility or any obligation whatsoever. The fact that the Government may have formulated or in any way supplied the said drawings, specifications, or other data, is not to be regarded by implication, or otherwise in any manner construed, as licensing the holder, or any other person or corporation; or as conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

The Office of Public Affairs has reviewed this report, and it is releasable to the National Technical Information Service, where it will be available to the general public, including foreign nationals.

This report has been reviewed and is approved for publication.



BYRON J. PIERCE
Project Scientist

Elizabeth L. Martin
ELIZABETH L. MARTIN
Technical Director


LYNN A. CARROLL, Colonel, USAF
Chief, Aircrew Training Research Division

Please notify AL/HRPP, 7909 Lindbergh Drive, Brooks AFB, TX 78235-5352, if your address changes, or if you no longer want to receive our technical reports. You may write or call the STINFO Office at DSN 240-3853 or commercial (210) 536-3853.

REPORT DOCUMENTATION PAGE

*Form Approved
OMB No. 0704-0188*

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED	
	June 1996	Final - June 1992 to May 1995	
4. TITLE AND SUBTITLE			5. FUNDING NUMBERS
Image Update Rate Can Affect the Perceived Speed of Simulated Self-Motion			C - F41624-95-C-5011 PE - 62205F PR - 1123 TA - B2 WU - 06
6. AUTHOR(S) Julie Mapes Lindholm Timothy M. Askins Norwood Sisson			8. PERFORMING ORGANIZATION REPORT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Hughes Training, Inc., Training Operations 6001 South Power Road, Bldg 561 Mesa, AZ 85206-0904			10. SPONSORING/MONITORING AGENCY REPORT NUMBER AL/HR-TR-1995-0194
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Armstrong Laboratory Human Resources Directorate Aircrew Training Research Division 6001 South Power Road, Bldg 558 Mesa, AZ 85206-0904			
11. SUPPLEMENTARY NOTES Armstrong Laboratory Technical Monitor: Dr Byron J. Pierce, (602) 988-6561			
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.		12b. DISTRIBUTION CODE	
13. ABSTRACT (<i>Maximum 200 words</i>) The update rate of an image generator affects the spatiotemporal representation of a time-varying scene and thus, potentially, an observer's percept during observation of the display image. In the first part of this report, we discuss image generation technology and sampling theory, and we present an analysis of the temporal frequencies in a space-time image representing constant-velocity, constant-altitude flight over a flat, textured terrain. In the second part, we report the results of two experiments in which a two-alternative, forced-choice method of constant stimuli was used to investigate the effects of image update rate (30 Hz vs 60 Hz) on the perceived speed of self motion. We found that perceived speed was higher with the lower update rate when the original image, internal to the computer, contained very high temporal frequencies, and the display image, therefore, contained a large number of spatiotemporal frequencies, within the bandpass of the human visual system, which had the wrong drift direction.			
14. SUBJECT TERMS Aliasing; Computer image generation; Flight simulation; Fourier transform; Image update rate; Motion perception; Optic array; Optic flow; Self motion; Spatiotemporal frequency; Speed perception			15. NUMBER OF PAGES 49
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION ABSTRACT SAR

CONTENTS

	<u>Page</u>
INTRODUCTION.....	1
Visual Simulation Technology.....	1
Sampling Theory.....	2
Spatial Frequencies in Images of Synthetic Environments.....	4
Temporal Frequencies in Images of Synthetic Environments.....	6
Constant Velocity Translation in a Three-Dimensional Environment.....	7
Spectral Consequences of Image Update Rate and Display Refresh Rate.....	10
Display Reconstruction Filters.....	13
Perceptual Significance of Discrepancies Between Display Image and Original Image.....	13
EXPERIMENT 1.....	15
Method.....	15
Observers.....	15
Apparatus.....	15
Spatial Images.....	15
Space-Time Images.....	17
Spatiotemporal Frequencies.....	19
Procedure and Design.....	23
Data Reduction and Analyses.....	25
Results and Discussion.....	25
EXPERIMENT 2.....	30
Method.....	30
Results and Discussion.....	32
GENERAL DISCUSSION.....	35
REFERENCES.....	39

List of Figures

<u>Figure</u>	<u>Page</u>
1 Fourier Transforms of a Continuous and a Sampled Image	3
2 Flow Pattern in Projection of Optic Array.....	8
3 Flat Ground Surface Textured with Sinusoidal Pattern.....	9
4 Spectral Effects of the Ratio of Refresh Rate to Update Rate.....	12
5 Texture Pattern.....	17
6 Effects of Magnification Filter on Texture Reconstruction.....	18
7 Spectrum of Texture Pattern.....	20
8 Slice of Spectrum of Texture Moving at a Constant Velocity.....	21
9 Attenuation of Spectral Components in Accord with Constant Velocity Translation with 30 Hz Update Rate.....	22
10 Slices of the Spectra of Spatially Local Regions of the Standard Display Images	24
11 Results of Experiment 1	27
12 Update-Rate Effect for Individual Observers.....	28
13 Texture Densities and Optic Flows for Experiment 2.....	31
14 Results of Experiment 2	33
15 Response Curves for Two Representative Observers.....	35

PREFACE

This research was conducted at the Armstrong Laboratory, Aircrew Training Research Division (AL/HRA), by University of Dayton Research Institute under Contract No. F33615-90-C-0005, Work Unit 1123-B2-06, Flying Training Research Support, and continued by Hughes Training, Inc., Training Division, under Contract No. F41624-95-C-5011. Ms. Patricia A. Spears and Mr. Dan Mudd were the Contract Monitors. Drs. Elizabeth L. Martin and Byron Pierce were the technical monitors.

We thank Ms Krystyna Krasnicka (AL/HRA) who assisted with the scheduling and testing, Ms Debra Bolin who helped prepare several of the figures, and the pilots and Laboratory personnel who served as subjects.

SUMMARY

The visual system of a typical flight simulator consists of an image generator (IG) and a display device. During a simulated flight, the IG samples, in both space and time, the time-varying perspective projection of the database-defined environment. The display system then creates a space-time image from the sequence of digital spatial images. The display image is never a perfect reconstruction of the scene that the pilot would see if he or she were able to actually fly through the synthetic environment. Many of the discrepancies between the display image and the scene it is meant to represent are due to the system's sampling and display processes.

In the first part of this report, we provide an overview of simulator technology and sampling theory. We go on to discuss the variables that determine the spatiotemporal-frequency spectrum of a time-varying image of a synthetic environment. Finally, we consider the effects of an IG's image-update (temporal sampling) rate and a display system's refresh rate and reconstruction filters on the spectrum of a display image.

The Fourier transform of a sampled version of an image is a periodic replication of the transform of the original image. The replicas along each dimension of the frequency space are separated by the sampling frequency for that dimension. Accurate reconstruction of an image from its sampled values requires that the spectral replicas occupy nonoverlapping regions of the frequency space and that during reconstruction the spectral replicas are removed.

With current image-generation technology, the reconstruction filters of the display device do not remove the replicas that result from temporal sampling. Moreover, the image that is sampled may contain temporal frequencies high enough to result in considerable spectral overlap. Regardless of the extent of spectral overlap, spatiotemporal frequencies introduced by sampling provide erroneous form and motion information. If they lie within the *window of visibility*, or bandpass, of the human visual system, they may affect perception.

The lower the image update (temporal sampling) rate of an IG, the more likely it is that portions of the spectral replicas will lie within the window of visibility. When the refresh rate of the display device is higher than the update rate of the IG, it will modify the spectrum of the image, but it will not compensate for the low update rate.

We conducted two experiments to examine the effects of visible, erroneous spatiotemporal frequencies on the perceived speed of constant-velocity, constant-altitude flight over a flat, textured ground surface. In this case, the temporal frequencies in the original image, internal to the computer, are the product of the spatial frequencies in the ground surface pattern and the simulated velocity, with both measured in environmental units. To manipulate the

visibility of the erroneous temporal-sampling-induced frequencies in the display image, we varied both the update rate of the IG (30 Hz vs. 60 Hz) and the temporal frequencies in the original image. The results of these experiments suggest that the perceived speed of self-motion is increased by the addition of a large number of visible, sinusoidal components with erroneous velocity information.

With current technology, there is a tradeoff between image-update rate and scene content. It is often considered desirable to reduce the update rate in order to maintain a relatively high level of detail. Our results indicate that for some combinations of ground speed and ground texture, reducing the update rate to 30 Hz will result in an increase in the perceived speed of self motion. This effect can be eliminated by lowering the rate at which a given texture pattern is traversed (i.e., by reducing ground speed or by mapping the texture pattern to a larger ground area) or by limiting the spatial frequencies in the texture pattern. Both of these solutions restrict the range of *temporal* frequencies in the original image.

IMAGE UPDATE RATE CAN AFFECT THE PERCEIVED SPEED OF SIMULATED SELF-MOTION

INTRODUCTION

Surfaces in an illuminated environment reflect light in all directions, creating a network of rays. As a consequence, light from all directions converges to each point from which an environment could be observed. The pattern of light reflected to any such point has been termed an *optic array* (Gibson, 1966). In general, there is a unique optic array at each convergence point, and the optic array at a *moving point of observation* (Gibson, 1979) changes over time. When an observer views an environment, each eye occupies an observation point and admits a portion of its optic array. This portion, at each point in time, is determined by the directions of the head and eye and the eye's field of view.

With computer image generation systems, it is possible to simulate the optic array for a moving point of observation in a synthetic, three-dimensional environment (Hochberg, 1986). This technology provides a means to investigate the visual determinants of self-motion perception and to train tasks (e.g., flying and driving) that are dependent upon the visual concomitants of movement through an environment. However, even the most advanced database-defined environments fail to match all of the potentially important characteristics of a "real-world" environment. Moreover, the spatiotemporal pattern of light produced by an image generation system is not a perfect representation of a portion of an optic array in the *synthetic* environment. Fidelity is reduced by imperfections in the system hardware as well as by certain inherent characteristics of the technology. Our concern here is with the perceptual consequences of a subset of the latter, more systematic distortions. To provide the background necessary to understand these distortions, we begin this report with an overview of simulator technology and sampling theory, followed by an analysis of the variables that determine the spatiotemporal-frequency spectra of computer-generated images. In this discussion, the synthetic environment is assumed to be monochromatic.

Visual Simulation Technology

In most flight and driving simulators, images are generated for a single point of observation (as if an observer had only one eye) and for a limited field of view (which may be either larger or smaller than an observer's instantaneous field of view). Except for some systems that monitor the position of the observer's head, the *projection reference point* (i.e., the designated point of observation in the synthetic environment) is fixed with respect to the simulated vehicle.

Within an image generator (IG), points of observation, optic arrays, view planes, and continuous images are mathematical abstractions. Bearing this in mind, the image-generation process can be conceptualized as follows: A digital database and a moving point of observation

determine a time-varying optic array. The optic array is intersected by a view plane, a windowed region of which defines the projection surface. As the vehicle moves through the synthetic environment, a continuous space-time image, $i(x,y,t)$, is formed on the window. This continuous image (or a spatially lowpass version of this continuous image) is sampled and quantized at a three-dimensional, spatiotemporal array of points, resulting in a sequence of digital, spatial images, each of which represents the optic array at a single point in time. The spatial images are successively loaded into the IG's graphics memory, where the pixel ("picture element") values are scanned (from left to right, from top to bottom) and converted into an analog signal that modulates the output of a display device, creating a visible, space-time image.

Sampling Theory

If a continuous three-dimensional image, $i(x,y,t)$, is multiplied by a sampling function consisting of a three-dimensional, regularly spaced array of impulses, a *sampled image* is obtained. The Fourier transform of this sampled image is a periodic replication of the Fourier transform of the original image, $I(u,v,w)$, where u is horizontal spatial frequency, v is vertical spatial frequency, and w is temporal frequency. For a sampling function with spacing Δx in the x direction, Δy in the y direction, and Δt in the t direction, the replicas in the spectrum of a sampled image are separated by $1/\Delta x$, $1/\Delta y$, and $1/\Delta t$ along the u , v , and w dimensions, respectively (see Gonzalez & Wintz, 1987; Jain, 1989; Stanley, Dougherty, & Dougherty, 1984, for discussions of sampling theory).

If the sampling frequency for each dimension is greater than twice the image's maximum frequency along that dimension, the spectral replicas will occupy nonoverlapping regions of the frequency space. The original image can then be recovered from the sampled image by a reconstruction filter that passes the spectrum of the original image and attenuates the sampling-induced spectral replicas. If an image is sampled at less than twice its maximum frequency for any dimension, then the baseband of the spectrum of the sampled image will contain frequencies *greater* than half that sampling frequency, and one or more of the sampling-induced spectral replicas will contain frequencies *less* than half that sampling frequency. This spectral overlap is known as *aliasing*, and when it occurs, the original function *cannot* be recovered by subjecting the sampled image to a lowpass reconstruction filter.

For example, assume that $i(y,t)$ is an image of a time-varying, one-dimensional spatial pattern and that the region of support of its Fourier transform (i.e., the region of the spatiotemporal-frequency space where the transform is different from zero) is as shown in the upper panel of Figure 1. To create a sampled version of this image, it is multiplied by a two-dimensional sampling function consisting of a train of impulses separated by Δy in the y direction and Δt in the t direction. If $1/\Delta y > 2v_{\max}$ and $1/\Delta t < 2w_{\max}$, then the support of the transform of the sampled image might look like that depicted in the lower panel of Figure 1. Because of the overlap of replicas along the w dimension, the original image cannot be reconstructed from the sampled values.





Figure 1
Fourier Transforms of a Continuous and a Sampled Image

The spectrum of a reconstructed image will depend upon the reconstruction filter. In the lower panel of Figure 1, the dotted rectangle specifies the region of the frequency space defined by $\pm 1/2\Delta y$ and $\pm 1/2\Delta t$. If the reconstruction filter passed all of the spatiotemporal-frequency components within this region and passed none of the components outside this region, the frequencies in the original image that exceeded half the sampling rate (i.e., those in regions a' and b') would be lost. In region a , the original nonzero components would be modified, and in region b , frequencies with nonzero amplitude would be added. Other reconstruction filters are possible. For example, with a reconstruction filter that passed everything below w_{\max} , the spectrum of the reconstructed image would include the uncorrupted components in b' as well as the corrupted components in a' .

If sampling frequencies are fixed and less than twice the highest frequency for one or more dimensions of an image, aliasing can be avoided (or at least substantially reduced) by application of a suitable lowpass filter *prior* to sampling. Appropriate filtering during reconstruction will then result in a lower resolution but otherwise distortion-free copy of the original image.

Although nonoverlapping replicas in the Fourier transform of a sampled image ensure that the original image *can* be recovered, they do not ensure that it *will* be recovered. If the reconstruction filter is too narrow, the reconstructed image will not include the higher frequencies of the original image; if the reconstruction filter is too wide, the reconstructed image will include frequencies from (one or more of) the sampling-induced spectral replicas.

Spatial Frequencies in Images of Synthetic Environments

In most visual simulation systems, environmental surfaces are represented by one or more polygons. Because polygon-defined surfaces have infinitely sharp edges, the continuous image on the view-plane window is generally not bandlimited. To reduce aliasing artifacts that result from polygon edges and from polygon projections that are very small in at least one direction, many IGs use supersampling and averaging techniques to filter the image prior to pixel sampling (Foley, van Dam, Feiner & Hughes, 1992; Watt, 1993). These antialiasing filters are typically applied *after* the intensity of each polygon has been computed for each pixel. As a consequence, although they limit the spatial frequencies of patterns formed *by* polygons, they have no effect on the spatial frequencies of patterns *on* polygons.

Patterned polygons can be created by systems that support texture mapping. A texture is stored in an IG as a digital array, usually of two dimensions. Its individual elements are called *texels*. Texel values can be generated either by executing a computational algorithm or by sampling an image. However, no matter what the actual source of the texel values, they should be taken to result from sampling a continuous spatial pattern at greater than twice its maximum horizontal and vertical spatial frequencies. Thus, a 256 x 256 array of texel values represents horizontal and vertical sampling frequencies of 256 samples/pattern and an original

spatial pattern with horizontal and vertical frequencies up to 128 cycles/pattern. In the sampled-pattern's periodic spectrum, replicas of the continuous-pattern's spectrum occur at multiples of 256 cycles/pattern along both dimensions of the frequency space.

During the image generation process, a texture is applied, or *mapped*, to a surface. Conceptually, texture mapping occurs in the three-dimensional environment, not in the two-dimensional image of that environment, and it includes reconstruction of a spatial pattern from the texel values. After a texture is mapped to a surface, it modulates the intensity of that surface and thus contributes to the optic array. Texture projections that contain horizontal or vertical spatial frequencies greater than half the pixel-sampling frequency¹ will result in aliasing unless that projection is subjected to a lowpass filter before it is sampled by the pixel array.

If a texture pattern was accurately reconstructed on an environmental surface and if the *projected spacing* between texels was consistently greater than the spacing between pixels (considering both texels and pixels as points), then the horizontal and vertical frequencies in the surface projection would be less than half the pixel sampling frequency, and aliasing would *not* occur. For a given textured surface, however, the projected spacing between texels typically varies over both space and time. In projections of relatively distant surface areas, texel spacing tends to be less, sometimes much less, than pixel spacing. Texel spacing may also be insufficient for relatively close surfaces that are very small (in environmental units) or that have a high slant.

Rather than using supersampling or related techniques to reduce aliasing of texture projections with high spatial frequencies, many simulation systems filter the texture *before* it is reconstructed on an environmental surface (see Heckbert, 1986, for a review of texture-antialiasing techniques). Several resolutions of each texture pattern are stored in the IG. (The lower the resolution, the fewer the texels and the lower the maximum possible spatial frequency in cycles per pattern.) During image generation, the IG determines, for each sampled point in time, the resolution necessary to prevent aliasing at each sampled point in space. It is this version (or an average of the two closest versions) that is reconstructed at the point on the environmental surface that projects to the point in the image that is sampled to obtain the pixel value. The version of a texture that is reconstructed (at points) on an environmental surface is thus determined independently for each spatiotemporal sample of the surface projection.

If texture-reconstruction and texture- and polygon-antialiasing filters were ideal, the spatial image sampled by the pixel array would correspond to a projection of the optic array for the prescribed synthetic environment, after that projection had been subjected to a lowpass filter that removed spatial frequencies greater than half the pixel-sampling frequency and that passed, without modification, frequencies less than half the pixel-sampling frequency. The spatial frequencies of near-surface projections would usually be limited by the texture resolution whereas those of far-surface projections would be limited by the pixel-sampling frequency. In

¹ An IG's horizontal sampling frequency and vertical sampling frequency are assumed to be equal in the following discussion.

practice, the original texture patterns are never perfectly reconstructed on environmental surfaces, and the image sampled by the pixel array is never perfectly bandlimited.

Temporal Frequencies in Images of Synthetic Environments

In the Fourier transform $I(u, v, w)$ of a space-time image, the points (u, v, w) and $(-u, -v, -w)$ specify a drifting sinusoidal grating with spatial frequency $f = (u^2 + v^2)^{1/2}$, orientation $\alpha = \arctan(v/u)$, and drift velocity $r = -w/f$ (see Graham, 1989). Although the direction of drift is necessarily parallel to the orientation of the component (i.e., orthogonal to the "bars" of the grating), the speed is unrestricted.

If an image is static over time, the region of support of its Fourier transform will lie on the uv plane at $w = 0$. If an image varies over time, in any way, its support will include some components with $w \neq 0$. The spectrum of an arbitrary time-varying image may contain any combination of components (with an amplitude > 0), including many with the same temporal frequency and different two-dimensional spatial frequencies as well as many with the same two-dimensional spatial frequency and different temporal frequencies.

A time-varying image of a synthetic environment will usually have infinitely high temporal frequencies (given that $w = -rf$), and if a spatiotemporal-frequency component with a temporal frequency greater than half the image-update (temporal-sampling) rate is not removed before the image is sampled, it will be aliased. Nonetheless, current real-time simulation systems do not have temporal-antialiasing capabilities (i.e., the pixel values for a given temporal sample are determined by the image at a single point in time).

Although IG's spatial antialiasing filters indirectly limit the temporal frequencies in the image that is sampled, they do *not* ensure that it is free of temporal frequencies greater than half the image-update rate. Consider, for example, a system with horizontal and vertical sampling frequencies of 30 cycles/deg (i.e., $\Delta x = \Delta y = 2$ arc min) and a temporal sampling frequency of 60 Hz ($\Delta t = 16.67$ msec): A spatial antialiasing filter that removed spatial frequencies greater than 15 cycles/deg would prevent temporal aliasing (i.e., temporal frequencies ≥ 30 Hz) if and only if the image velocities were all less than 2 deg/sec.

High image velocities are most likely in projections of near surfaces. The spatial frequencies of near-surface projections may, however, be limited by the texture pattern rather than by the pixel-sampling frequency. (As previously noted, if the projected spacings between texels are greater than the spacing between pixels, the maximum horizontal and vertical frequencies in the projection of a surface--upon which a texture pattern has been accurately reconstructed--will be less than half the pixel sampling frequency.) In this case, the spatial frequencies in the surface projection may be much lower than half the pixel sampling rate, and temporal aliasing may not occur even for very high image velocities. For example, if the maximum, texture-limited spatial frequency in the projection of a near surface was only 0.5

cycles/deg, the local image velocity would have to exceed 60 deg/sec to cause temporal aliasing on a system with a 60 Hz update date.

Constant Velocity Translation in a Three-Dimensional Environment

As we have discussed, many characteristics of a synthetic environment and of the point of observation's motion within that environment affect the spatiotemporal-temporal frequencies in the original space-time image. Moreover, prior to sampling, the temporal frequency spectrum as well as the spatial frequency spectrum of the original image is usually modified by the IG's spatial-antialiasing filters. In the research reported here, we restrict our attention to the images that result during constant velocity, constant altitude translation of a point of observation in a synthetic environment consisting of a flat, textured ground surface and an unpatterned sky.

The spatiotemporal frequencies in such images can be understood by first considering the simpler case of constant velocity translation in the image plane. As described fully by Watson and Ahumada (1985), if an image consists of a two-dimensional spatial pattern translating at a constant velocity, the region of support of its Fourier transform lies in an oblique plane through the origin, with the orientation of the plane uniquely defining the image velocity. The temporal frequency of each point in the spectrum equals minus the dot product of the image-velocity vector $\mathbf{r} = (r_x, r_y)$ and the spatial-frequency vector $\mathbf{f} = (u, v)$, that is, $-(r_x u + r_y v)$.

If a view-plane window maintains a constant orientation and position in front of a point of observation that is moving forward at a constant velocity and altitude over a flat surface, the instantaneous velocities in the image vary across spatial positions, as depicted in Figure 2 (Gibson, 1950). In addition, the projection of a given surface feature depends upon its position with respect to the point of observation (see Fig. 3). Clearly then, constant-velocity translation over flat terrain does not result in an image that translates at a constant velocity. Nor can such images be decomposed into spatiotemporally local regions of constant-velocity translation. Although their temporal-frequency composition is dependent upon a single velocity, the critical velocity is not that of the image but that of the ground plane with respect to the point of observation: The temporal frequencies in the image spectrum will equal minus the dot products of the two-dimensional spatial-frequency vectors in the surface spectrum (in cycles per environmental unit or per texture pattern) and the surface-velocity vector (in environmental units or texture patterns per second).

Figure 3 illustrates computer-generated spatial images of three views of a flat terrain that is textured with a single sinusoidal grating; the altitude of the point of observation equals 2.5 times the period of the grating. Clearly, for these viewing conditions, the projection of a *single* spatial frequency is a pattern that contains *many* different spatial frequencies. Although the spatial-frequency components in the top panel of Figure 3 all have a vertical orientation (i.e., the v axis is the support of the Fourier transform), the components in the other two images differ in both orientation and frequency.




Figure 2
Flow Pattern in Projection of Optic Array

Regardless of the viewing direction, in an image formed during forward motion over a flat surface textured with a sinusoidal grating, the instantaneous velocity of a grating "bar" at a given point along its contour will equal the component of the point velocity (Fig. 2) that is orthogonal to the contour. The region of support of the Fourier transform of such a space-time image will lie in planes at minus and plus a single temporal frequency.

Consider, for example, the original, ideal (with perfect reconstruction of the sinusoidal pattern on the environmental surface) space-time images that are formed if the spatial frequency of the texture pattern in Figure 3 is five cycles per environmental unit and the point of observation moves forward (and thus the ground surface moves backward) at three environmental units per second. For the view depicted in the top panel of Figure 3, the motion in the three-dimensional environment is orthogonal to the bars of the grating. In the image, each "bar" of the grating will move down, with proportionate increases in size and speed. The luminance at each point in the ground-surface projection (except those on the horizon) will modulate sinusoidally at 15 Hz, and in the spectrum of the image, the two points defining each (non-zero) spatiotemporal-frequency component will have temporal frequencies of plus and minus 15 Hz--minus the dot product of the surface-velocity vector, $\mathbf{r} = (0, -3)$, and the spatial-frequency vectors of the spectrum of surface pattern, $\mathbf{f}_1 = (0, 5)$ and $\mathbf{f}_2 = (0, -5)$. For the view depicted in the middle panel of Figure 3, the point of observation moves parallel to the bars of the grating. In the image, therefore, the spatial pattern will be stationary, and in the spectrum of the image, each spatiotemporal-frequency component will have a temporal frequency of 0 Hz ($-\mathbf{r} \cdot \mathbf{f} = 0$). For the view depicted in the lower panel of Figure 3, the point of observation moves at a 45-degree angle with respect to the grating. In the image, the grating bars will move down and to the left, with both the speed and direction of motion varying horizontally as well as vertically. The luminance at each point in the ground-surface projection will modulate sinusoidally at a frequency of 10.6 Hz, and in the spectrum of the image, the two points defining each component




Figure 3
Flat Ground Surface Textured with Sinusoidal Pattern

will have temporal frequencies of plus and minus 10.6 Hz, that is, $-\mathbf{r} \cdot \mathbf{f}$, where $\mathbf{f}_1 = (3.5, 3.5)$ and $\mathbf{f}_2 = (-3.5, -3.5)$

Spectral Consequences of Image Update Rate and Display Refresh Rate

Thus, for constant-velocity, constant-altitude translation over a flat, textured terrain, the spatial pattern of the terrain (after any spatial antialiasing filters have been applied) and the altitude and velocity of the point of observation determine the spatiotemporal frequencies in the image that is sampled. Both the update rate of the IG and the refresh rate of the display device affect the frequency composition of the resulting display image.

The image-update (temporal-sampling) rate determines the spacing of the spectral replicas along the temporal-frequency dimension of the Fourier transform of a sampled image. However, the refresh (vertical scan) rate of the display device may be higher than the update rate of the IG, so that each digital spatial image is displayed k times, where $k:1$ is the ratio of the display's refresh rate to the IG's update rate. (Although the update rate of some IGs is allowed to vary over time, for the purposes of this discussion the update rate of a given IG will be assumed to be constant.) For $k > 1$, the spatiotemporal array (i.e., the sampled-image) that is scanned by the display system will not match the spatiotemporal array that results from sampling the original image. In this case, the display image is a "reconstruction" of an image that differs from the original. For example, if the original image is for a point of observation moving through an environment at a constant velocity \mathbf{r} , the display image is a reconstruction of an image (sampled at the display rate) for a point of observation that alternates between $k - 1$ periods of no motion and 1 period of motion with velocity $k\mathbf{r}$.

The typical derivation of the periodic Fourier transform of a time-sampled function involves noting that multiplication in the time domain is equivalent to convolution in the frequency domain, and that the Fourier transform of a train of impulses with spacing Δt is a train of impulses with spacing $1/\Delta t$. In the case of constant-velocity translation of an image, an alternative derivation is possible: First, a space-time image consisting of a "single presentation" of the spatial pattern is constructed by multiplying the spatial image by a temporal impulse. Then, to construct a temporally sampled version of the image moving at velocity \mathbf{r} , the single-presentation image is convolved with a sampling function in which successive impulses are separated by both Δt in time and $\mathbf{r}\Delta t$ in space. Because convolution in the space-time domain is equivalent to multiplication in the spatiotemporal-frequency domain, the spectrum of this sampled image is the product of the spectrum of the impulse train and the spectrum of the single presentation of the spatial pattern.

The advantage of this alternative derivation is that it can easily be modified to show the spectral consequences of an update rate that is less than the refresh rate. For example, if $k = 2$ the spatial image is multiplied by a pair of impulses (rather than by one), separated by the display refresh interval, Δt_D , thereby constructing a space-time image consisting of two presentations of

the spatial pattern. Convolution of this image with a sampling function in which successive impulses are separated by Δt_U in time and by $r\Delta t_U$ in space, where Δt_U is the update interval of the IG, results in a sampled version of an image in which a spatial form is alternately stationary for Δt_D and translating at $2r$ for Δt_D . The spectrum of this sampled image is the product of the spectrum of the impulse train and the spectrum of two presentations of the spatial pattern.

Figures 4a to 4f illustrate this derivation, for $k = 1$ (left column) and $k = 2$ (right column), for an original image consisting of a horizontal bar (of height $4\Delta y$) moving at a velocity $r = -\Delta y/\Delta t$. For $k = 2$, $\Delta t_U = 2\Delta t$ and $\Delta t_D = \Delta t$, where $1/\Delta t$ is both the update rate and refresh rate for $k = 1$. In each figure, the horizontal dimension is temporal frequency, the vertical dimension is vertical spatial frequency, and the brightness of a point indicates its magnitude. The depicted regions of the spectra extend from $-1/2\Delta t$ to $1/2\Delta t$ along the temporal frequency dimension and from $-1/2\Delta y$ to $1/2\Delta y$ along the vertical frequency dimension; the actual spectra would extend from $-\infty$ to ∞ along both dimensions.

In the spectrum of an image consisting of a single presentation of the bar (Fig. 4a), the amplitude of each spatial frequency is constant over the full range of temporal frequencies, whereas in the spectrum of an image of two presentations of the bar (Fig. 4b), the amplitude varies according to $\cos \pi s$, where $s = w\Delta t_D$ (Bracewell, 1986), with the first zero at half the refresh rate (i.e., at the update rate). Multiplication of the Fourier transforms of the single-presentation and dual-presentation images by the transforms of the appropriate sampling functions (parallel line impulses with slope $-1/r$ and spacing $1/\Delta t$ [Fig. 4c] and $1/\Delta t_U$ [Fig. 4d], respectively) results in the transform of a sampled version of an image in which the bar is moving at a velocity r (Fig. 4e) and the transform of a sampled version of an image in which the bar is alternately stationary for Δt_D and moving at a velocity $2r$ for Δt_D (Fig. 4f). Within the depicted region of the frequency space, the support of the spectrum for $k = 1$ is a line through the origin with a slope of $-1/r$, whereas the support of the spectrum for $k = 2$ includes both this constant velocity line and parallel lines introduced by the lower update rate. Every spatiotemporal-frequency component that falls on these additional lines represents a sinusoidal grating that is drifting in the “wrong” direction. (Each point in the second quadrant can be connected to its corresponding point in the fourth quadrant by a line through the origin with a slope of $-1/r_i$, where r_i is the velocity of the component.)

The spectra in Figures 4e and 4f also differ in the amplitudes of components falling on the constant-velocity line. For $k = 1$, they correspond to the amplitudes of the spectrum of the original spatial form (i.e., a horizontal bar with height $4\Delta y$); for $k = 2$, they decrease according to $\cos \pi s$. For $k = 2$ then, the spatiotemporal-frequency components with a drift velocity of r define a form that differs from that in the original image (Lindholm, 1992a). More generally, as inspection of Figure 4b indicates, for $k = 2$ the spatial spectrum of a constant velocity line will be a function of that velocity.




Figure 4
Spectral Effects of the Ratio of Refresh Rate to Update Rate

Finally, it is important to consider the spectrum that would result if the display refresh rate, $1/\Delta t_D$, matched the lower update rate, $1/\Delta t_U$ --that is, if the image were both sampled and displayed at $1/2\Delta t$. For the region of the frequency space displayed in Figure 4, the support of this spectrum would correspond to that depicted in Figure 4f, but the spatial spectrum of the constant velocity line would match that in Figure 4e. The "extra" parallel lines, in this case, would represent the upper and lower halves, respectively, of the spectral replicas located at $-1/2\Delta t$ and $1/2\Delta t$.

Display Reconstruction Filters

As a simulator's display system scans the sampled image (i.e., the sequence of digital, spatial images), it reconstructs a space-time image from the sampled values. The spatial reconstruction filtering, which involves a sample and hold process and the intensity distribution of the CRT spot, corresponds to convolving the sampled values first with a 1-pixel wide box filter and then with a Gaussian (Foley et al., 1992). Although this composite filter passes some frequencies greater than half the pixel sampling frequency and attenuates some frequencies less than half the sampling frequency, it can be considered to be an adequate approximation of the ideal (i.e., to remove the spatial-sampling-induced spectral replicas while passing the spatial frequencies in the baseband spectrum). In contrast, the temporal filtering of the display device is limited to that provided by the temporal impulse response of the phosphor. Most phosphors used in simulator displays decay to 10% of their maximum intensity in less than a millisecond. The *temporal*-sampling-induced spectral replicas are thus passed with *little* attenuation.

Perceptual Significance of Discrepancies Between Display Image and Original Image

Although a display image is never a perfect reconstruction of the original image, some discrepancies between the two will involve spatiotemporal frequencies that lie outside the *window of visibility* (Watson, Ahumada, & Farrell, 1986)--that is, outside the spatiotemporal-frequency passband of the human visual system (during fixation of a stationary mark). According to many models of the human visual system, such discrepancies will not affect perception: The human visual system functions as a final filter between a display image and visual perception. The display image and the original image will thus be perceptually equivalent if their spectra, after passing through the window of visibility, do not differ.²

Discrepancies between original and display images can be divided into two categories, those that involve a loss of information and those that involve erroneous information. The former arise when the presampling (antialiasing) filters remove or attenuate components in the

² Although we will disregard the effects of eye movements in this discussion, it is important to note that if the direction of gaze varies during the observation of an image, the temporal frequencies in the *retinal* image can differ substantially from those in the display image. Differences between the display image and the original image that would not be visible if the direction of gaze were fixed may be visible if smooth pursuit eye movements occur, and vice versa (Hsu, 1985; Lindholm, 1992a).

original image or when the display-reconstruction filters remove or attenuate components in the baseband of the sampled-image spectrum. The latter arise when the sampling process corrupts or introduces components that are then passed by the reconstruction filters or when the system's output pattern distorts the spectrum.

With current technology, lost information typically involves spatiotemporal-frequency components with high spatial frequencies. If these frequencies fall within the passband of the human visual system, the display image will usually appear blurred, and form detection and discrimination may be more difficult than they would have been if the high spatial frequencies had not been lost.

Erroneous information often results from high temporal frequencies in the original image or from high temporal frequencies introduced by temporal sampling and passed by the temporal reconstruction filter. If temporal frequencies greater than half the update occur in the image that is sampled, the temporal-sampling-induced spectral replicas will overlap. In the display image, the resulting aliased spatiotemporal frequencies will have temporal frequencies less than half the update rate. They will often fall within the window of visibility. As demonstrated by the wagon-wheel effect and related phenomena, temporal aliasing can result in misperception of the velocity (speed and direction) of a cyclic spatial pattern.

Because the temporal-sampling-induced spectral replicas are not removed by the display device during reconstruction, high temporal frequencies in the original image will also appear in the display image. (As illustrated in the lower panel of Fig. 1, the spatiotemporal frequencies in this portion of the baseband spectrum may [region a'] or may not [region b'] be modified.) If the update rate is twice the maximum temporal frequency of the window of visibility, these components, as well as components in the spectral replicas that exceed half the sampling rate, will fall outside of the window of visibility. Otherwise, some of the components in the spectral replicas will fall within that window and may affect perception. Note in particular that an absence of spectral overlap does not ensure that the window of visibility will be free of spatiotemporal frequencies introduced by temporal sampling.³ (Consider, for example, the flicker observed in a "static" display image.)

The display image will also contain erroneous information when the update rate of the IG is lower than the refresh rate of the display device. As we have discussed, the image that is scanned during reconstruction does not, in this case, correspond to the one that is sampled, and many of the differences between them can be expected to fall within the window of visibility. Prior research has shown that the ratio of refresh rate to update rate can affect the perceived form

³ The term aliasing has been reserved, thus far, for the spectral overlap that results from an inadequate sampling rate. More generally, every waveform that fits a given set of sampled values is considered an alias (Stanley et al., 1984). With this definition, each component in a spectral replica is an alias of the corresponding component in the baseband spectrum. Regardless of spectral overlap, then, display images are always characterized by aliased frequencies introduced by temporal sampling.

of a moving object (Hempstead, 1966; Stenger, Zimmerlin, Thomas, & Braunstein, 1981; Lindholm, 1992a, 1992b).

In the experiments reported here, we investigated the effects of erroneous spatiotemporal frequencies on the perceived speed of self-motion over a textured, planar surface. We limited our consideration to the case in which the point of observation moves forward at a constant velocity and altitude. The amount of erroneous information was manipulated by varying both the update rate of the IG and the spatiotemporal frequency content of the original image. In the first experiment, the speed of the point of observation was varied. In the second experiment, texture rate and optic flow rate were dissociated by varying the optical texture density.

EXPERIMENT 1

Method

Observers. Eight male United States Air Force pilots served as observers. All had normal or corrected-to-normal vision and had both low-level flight and simulator experience. They were informed that we were investigating the effects of image update rate on velocity perception.

Apparatus. We used a Silicon Graphics SkyWriter to generate the motion sequences (space-time images). The Skywriter provides 24-bit color, double buffering, correctly interpolated perspective projections, and real-time texture mapping with filters for reconstruction and antialiasing. It is equipped with two complete graphics subsystems (pipelines). In order to achieve a reliable 60 Hz update rate and a spatial resolution of 1280 Horizontal (H) x 1024 Vertical (V), the two pipelines were programmed to generate alternate spatial images, and priorities were set to prevent the UNIX operating system from interrupting a motion sequence.

The analog output of the Skywriter was sent to a Barco 800 CRT Projector. The 60 Hz, noninterlaced output of the Barco was focused on a 190 x 150 cm area of a 240 x 180 cm, uniform-gain, rear-projection screen. Prior to each test session, projector convergence was checked and, if necessary, adjusted. Then, to ensure that the output levels of the CRTs were not varying systematically over the course of the experiment, a Minolta TV-2150 Color Analyzer was used to measure the display's chromaticity and luminance for the digital values representing black, white, and the experimental ground and sky colors.

Spatial Images. Each spatial image represented a sector of the optic array, at an instant in time, for a point of observation in a synthetic environment. The view-plane normal was parallel to a line from the point of observation to the center of the view-plane window. The window's up-axis was perpendicular to the ground surface. The vertical view angle was set to 60 degrees. Given the 1.25 aspect ratio of the rectangular view-plane window, the horizontal view angle was about 72 degrees.

The synthetic environment consisted of a flat, textured terrain and an untextured sky. The terrain was effectively composed of an infinite string of identical, abutting squares, extending in depth; however, near and far clipping planes limited the visible "length" of the ground surface to six squares.

The altitude of the point of observation was 1/120 of the length of a terrain-square. With six of these squares visible, the "terrestrial horizon" (Sedgwick, 1986) formed by the far boundary of the rectangular ground surface was less than five arc minutes below the true horizon. The center of a black fixation cross, formed by two 21-pixel lines and subtending about 1.2 deg x 1.2 deg, was positioned approximately 1.3 deg below the true horizon.

The color of the sky was set to a uniform blue (RGB = 0, 112, 192); the terrain squares were assigned a uniform brown (RGB = 255, 160, 96). The actual RGB values for a terrain pixel were computed by multiplying *each* of the assigned values by a filtered texel value (between 0 and 1). Before being converted to an analog signal, the digital pixel values were passed through the Skywriter's "gamma-correction" look-up table, which was set with the default value of 1.7. (This value would linearize the output of the system if the function relating digital value and output luminance were related by a gamma function with an exponent of 1.7). The Minolta Color Analyzer did not provide unbiased estimates of luminance or chromaticity. However, the measured luminance of the assigned ground color was roughly double that of the assigned sky color. Our best estimate of the former was about 10 cd/m².

The texture pattern is illustrated in Figure 5. To create this pattern, we (a) generated a 256 x 256 array of randomly selected values between 0 and 255, (b) subjected this digital image to a lowpass filter, and (c) expanded the restricted range of values produced by the filtering process. A histogram of the final texel values approximated a normal distribution, with values ranging from 0 to 255.

During image generation, the texel value for a displayable pixel was obtained by application of either a *bilinear magnification filter* or a *MIPmap trilinear minification filter* to a local area of the texture pattern. The bilinear magnification filter was used to reconstruct the texture on the near terrain, where the projected texel spacing was greater than the pixel spacing. This filter computed a weighted average of the values of the four texels nearest to the point in the texture to which the pixel mapped. The MIPmap trilinear minification filter was used to reconstruct the texture and to minimize aliasing on the far terrain. This filter selected the two most appropriate texture resolutions (from a stored set of prefiltered versions in which each successive version had half the resolution of the preceding version) and then computed a weighted average of the values of the four texels, in each of the selected versions of the texture, nearest to the point in the texture to which the pixel mapped.




Figure 5
Texture Pattern

The effects of the magnification filter are illustrated in Figure 6. The image in the top panel depicts the terrain when "point-sampling" was used to reconstruct the texture at those points where the projected texel spacing was greater than the pixel spacing. This technique is equivalent to convolving the texel values with a 1-texel-wide box filter (in accord with the common conception of a texel as a square of constant intensity). Note the "boxy" quality of the texture pattern and the sharp, jagged texel boundaries. In contrast, as depicted in the bottom panel, use of the bilinear magnification filter for reconstruction resulted in a texture pattern characterized by gradual changes in value and a rather "natural" appearance. The reconstructed texture was not, however, completely free of perceptible artifacts introduced by texel sampling: Some regions of the near terrain contained bright and dark lines parallel to the texel "edges."

Space-Time Images. Each space-time image represented a sector of the optic array as the point of observation moved forward at a constant velocity and altitude for a period of one second. The view-plane window was in front of the point of observation and perpendicular to the direction of motion. The IG's update rate was either 30 Hz or 60 Hz, whereas the display's refresh rate was always 60 Hz, noninterlaced. The point of observation moved parallel to the





Figure 6
Effects of Magnification Filter on Texture Reconstruction

longer side of the rectangular ground surface, which coincided with the vertical axis of the terrain texture.

Three standard (St) speeds were used: 0.1, 0.2, and 0.4 texture squares per second (ts/sec). A set of seven comparison (Co) speeds was associated with each St speed. Within a set, the log relative Co speed (i.e., $\log [Co/St]$) ranged from -0.15 to 0.15, in steps of 0.05. With a 60 Hz update rate, the three St speeds corresponded to displacements of 0.43, 0.85, and 1.71 texels per update; with a 30 Hz update rate, they corresponded to displacements of 0.85, 1.71, and 3.41 texels per update.

Spatiotemporal Frequencies. The amplitude spectrum of the texture's Fourier transform is depicted as a halftone, gray-scale image in Figure 7.⁴ The horizontal dimension, u , extends from $-1/\Delta x$ to $1/\Delta x$; the vertical dimension, v , extends from $-1/\Delta y$ to $1/\Delta y$. This spectrum should be understood to represent only a rough approximation of the spectrum in which we are interested, namely that of the pattern reconstructed on the ground surface adjusted for any nonlinearity in the function relating pixel value and display luminance.

With this limitation in mind, Figure 7 suggests that, for the higher values of v , the amplitudes of the two-dimensional spatial frequencies with $u = 0$ were as high as or higher than the amplitudes for any other value of u . Moreover, because the direction of motion in the three-dimensional environment corresponded to the vertical dimension of the texture, the horizontal frequency u had no effect upon the temporal frequency w of any spatiotemporal-frequency component in an image. Consequently, we will use the $(0,v)$ frequencies in the texture to illustrate the effects of speed and update rate on the image spectra.

The amplitudes of the $(0,v)$ frequencies are plotted in Figure 8. Their temporal frequencies in the original images for the St velocities (of the ground surface relative to the point of observation) are indicated by the three lower scales. Note that the maximum temporal frequencies were 12.8 Hz, 25.6 Hz, and 51.2 Hz for speeds of 0.1 ts/sec, 0.2 ts/sec, and 0.4 ts/sec, respectively.

⁴ When viewing this and subsequent gray-scale representations of two-dimensional spectra, it is important to be aware that the function used to map amplitudes to digital image values has a substantial effect upon the spectrum's appearance. Because many image spectra decrease rapidly with increasing frequency, only the very lowest frequencies may be visible if the mapping is linear. If $\log (1 + |I(u,v)|)$ values are mapped to the digital gray levels, the image depends upon the scale factor for the FFT. Whatever the mapping, the spectrum's appearance is also affected by nonlinearities in the response of the output device to the digital values and in the response of the human visual system to the pattern created by that device. For this report, the digital-image representations of Fourier spectra (except those in Figure 2) were generated by the following series of operations: Each amplitude value was divided by the maximum (dc) amplitude, any scaled values less than 10^{-4} were set equal to 10^{-4} ; the truncated set of scaled values was subjected to a log transformation, and the log values were linearly mapped to the 256 digital image values. The "gamma correction" (transfer function) for the printer was set so that the appearance of the printed output approximated that of the true gray-scale image on the computer monitor.






Figure 7
Spectrum of Texture Pattern

With a 60 Hz update rate, the replicas in the spectrum of the sampled image will not overlap if the temporal frequencies in the original image are all less than 30 Hz. For the St speeds of 0.1 ts/sec and 0.2 ts/sec, this condition was met. For the St speed of 0.4 ts/sec, spatiotemporal frequencies with temporal frequencies between 30 and 51.2 Hz were aliased, creating components in the image with temporal frequencies between 30 and 8.8 Hz. As illustrated in Figures 7 and 8, however, most of these components were of very low amplitude (note the log amplitude scale).


When the update rate was 30 Hz, the refresh rate was 60 Hz. Each digital spatial image was therefore displayed twice. As we discussed in the Introduction, for $k > 1$, where k is the number of times each digital spatial image is displayed, the sampled image that is scanned during reconstruction is not a sampled version of the original image. If $k = 2$ and the original image consists of a spatial pattern translating at a velocity \mathbf{r} , then the image that is scanned represents a sampled version of an image in which the spatial pattern is alternately stationary for Δt_D and moving at a velocity $2\mathbf{r}$ for Δt_D . In the spectrum of the sampled image, the amplitudes of the spatial frequencies that fall on the constant-velocity plane do not match those of the spatial




$r = (0.0, -0.1)$:



$r = (0.0, -0.2)$:



$r = (0.0, -0.4)$:



Temporal Frequency (Hz)

Figure 8
Slice of Spectrum of Texture Moving at a Constant Velocity

pattern, but are attenuated by $\cos \pi s$, where $s = w\Delta t_D$ (see Fig. 4). This also holds for images formed during constant velocity translation of a point of observation over a flat ground surface. Figure 9 illustrates the attenuation of the spatiotemporal-frequency components consistent with constant-velocity translation for each of the three St velocities.

In addition to these components, the sampled images that were scanned when the update rate was 30 Hz contained bands of energy centered at 30 Hz. These bands were parallel to but not replicas of the spectrum consistent with constant velocity translation. Their spectra corresponded to that in Figure 7, shifted by 30 Hz, and multiplied by $\cos \pi s$. Thus, the dc




Figure 9

Attenuation of Spectral Components in Accord with
Constant Velocity Translation with 30 Hz Update Rate

component of the original spectrum was reduced to zero in the spectral band at 30 Hz. The attenuation decreased as the temporal frequency of the components decreased.

Although we do not have a measure of the passband of the human visual system for this display, we assume that it extended to about 30 Hz (Watson et al., 1986). Thus, regardless of the maximum temporal frequency in the original image, portions of the bands of energy that were introduced by the 30 Hz update rate were likely to have fallen within the window of visibility and to have provided erroneous velocity information.

Figure 10 illustrates two-dimensional “slices” of the spectra of spatially local regions of space-time display images for the three St velocities (upper row: 60 Hz; lower row: 30 Hz; left to right: 0.1 ts/sec, 0.2 ts/sec, and 0.4 ts/sec). To create these spectra, we positioned a 128 x 128 pixel window on the horizontal center and 128 pixels below the vertical center of a display image. (The top edge fell just below the image position where the MIPmap minification filter was applied.) We then (a) saved the pixel values within this spatial window for each frame of a 60-sec motion sequence, (b) padded each end of the temporal sequence with two 128 x 128 spatial images in which every pixel was assigned a value of 128 (representing the homogeneous display between motion sequences), (c) multiplied the 128 x 128 x 64 array by a three-dimensional Gaussian window, (d) subjected the windowed values to a three-dimensional Fourier transform, and (e) took a slice at $u = 0$ through the resulting spectral volume.

The wv planes in Figure 10 extend from -30 Hz to 30 Hz horizontally and from -1/2 cycle/pixel to 1/2 cycle/pixel (± 8.5 cycles/deg) vertically. Points in the first and third quadrants (top right and bottom left) indicate components whose motion is consistent with the motion of the ground plane. Points in the other two quadrants represent components drifting in the wrong direction. Note that each of the anticipated temporal frequencies occurs with numerous spatial frequencies as a result of the perspective projection of the texture component with that temporal frequency (see Fig. 2). In general, the temporal frequencies in these spectra are completely consistent with those in Figures 8 and 9, except for a few low-amplitude components with higher-than-expected temporal frequencies. Because these components have relatively high spatial frequencies and velocities that are consistent with the ground velocity, they probably represent the high spatial frequency, texture-reconstruction errors that were visible in the spatial image.

Procedure and Design. The observers were tested individually for four sessions, each of which lasted 1.0 to 1.5 hours. The heights of the observer's chair and a chin rest were adjusted at the beginning of each session so that the observer's eyes were level with the vertical center of the display. From the viewing distance of 130 cm, the display image subtended approximately 72 x 60 deg of visual angle, in accord with the view angle specified by the IG. The images were viewed binocularly with natural pupils in a room with minimal ambient lighting.




Figure 10
Slices of the Spectra of Spatially Local Regions
of the Standard Display Images

Two 1-sec motion sequences were presented on each trial. One was selected from the set of 3 St speeds (0.1, 0.2, 0.4 ts/sec); the other was selected from the set of 7 Co speeds ($\log [Co/St] = -0.15$ to 0.15, in steps of .05) associated with the selected St. Co update rate (60, 30 Hz) and St update rate (60, 30 Hz) were varied independently to create 4 Co and St update-rate combinations: 60 & 60, 30 & 30, 60 & 30, and 30 & 60 Hz, for the Co and St, respectively. The 2 motion sequences for a given trial were separated by a 1-sec homogeneous field of the assigned sky color plus the black fixation cross. A 2-sec presentation of the same image occurred between the observer's response and the next trial. Observers were instructed to maintain fixation on the black cross and press the left button of a computer mouse if the first "flight" seemed faster and press the right button of the computer mouse if the second flight seemed faster.

At the beginning of the first session, an extended motion sequence was presented while the general nature of the experiment was explained. Observers were then presented with 24 practice trials consisting of the 2 easiest discriminations ($\log Co/St = \pm 0.15$) for each St, at 60 & 60 and 30 & 30, with the St presented first on half of the trials and the Co presented first on the other half of the trials. After the practice trials, observers were allowed to ask questions and take a short break. Before beginning the test sessions, they were instructed that many of the speeds would be very hard to discriminate and that it was important they always give their "best guess."

Over the 4 test sessions, each observer viewed each of the 84 (3 St speeds x 7 Co speeds x 4 update-rate combinations) motion-sequence pairs a total of 16 times. Each session was divided into 4 blocks of 84 trials. Two block types were presented, one containing the 60 & 60 and 30 & 30 update-rate combinations and the other containing the 60 & 30 and 30 & 60 combinations. Each of the two block types was presented twice each session, in alternate order. The type presented first was counterbalanced across subjects and sessions. Within a block, each of 42 motion-sequence pairs (3 St speeds x 7 Co speeds x 2 update-rate combinations) was presented twice, once with the Co first and once with the St first. The trial sequence was randomized independently for each observer.

Data Reduction and Analyses. For each observer, the proportion of trials on which the Co motion sequence was judged to be faster than the St motion sequence was computed for each of the 84 motion-sequence pairs. These proportions of "Co-Faster" judgments were subjected to a repeated measures analysis of variance (ANOVA), with 4 update-rate combinations, 3 St speeds, 7 $\log(Co/St)$ values, and 8 observers. Specific effects were assessed by performing ANOVAs on the appropriate subset of data. In addition, probit analysis (Finney, 1971) was used to obtain estimates of the points of subjective equality (PSEs) for a subset of the psychometric functions.

Results and Discussion

The update-rate-combination and $\log (Co/St)$ main effects and the update-rate-combination x $\log (Co/St)$ interaction were all statistically significant, $F_{3,21} = 21.48, p < .0001$ or

$p < .0006$ with the Greenhouse-Geisser (G-G) adjustment; $F_{6,42} = 231.65, p <.0001$; and $F_{18,126} = 3.58, p <.0001$ or G-G $p < .02$, respectively. Although the significance of the main effect of St speed was equivocal, $F_{2,14} = 4.55, p < .03$ or G-G $p < .06$, the effect of update-rate combination varied significantly across St speeds, $F_{6,42} = 23.53, p < .0001$, as did the interaction of update-rate combination and log(Co/St), $F_{36,252} = 2.56, p < .0001$ or G-G $p < .05$.

The three-way interaction of update-rate combination, log (Co/St), and St speed is depicted in Figure 11, where the mean percentage of "Co-Faster" responses is plotted as a function of log (Co/St) for each update-rate combination, for each St speed. Inspection of the data suggests that the update-rate combination had little, if any, effect upon the psychometric functions for St speeds of 0.1 ts/sec (left panel) and 0.2 ts/sec (middle panel). This observation was supported by subsequent ANOVAs. For the 0.1 ts/sec St, neither the main effect of update-rate combination nor the interaction of update-rate combination and log (Co/St) approached significance ($F < 1$, for both effects). Averaged over the four update-rate combinations, as log (Co/St) increased from -0.15 to 0.15, the mean percent of "Co-Faster" responses increased regularly from about 14% to about 90%, $F_{6,42} = 72.84, p < .0001$. For the 0.2 ts/sec St (Fig. 11, middle panel), the main effect of update rate combination was of borderline significance, $F_{3,21} = 3.29, p < .05$, G-G $p < .07$, with the highest mean percent (53.2%) for the 30 & 60 combination and the lowest mean percent (48.5%) for the 60 & 30 combination. The update-rate combination \times log (Co/St) interaction was not significant ($p > .05$, G-G $p > .20$); the mean percent of "Co-Faster" responses increased from about 9% to about 95% as log (Co/St) increased from -0.15 to 0.15, $F_{18,126} = 158.91, p < .0001$.

In contrast to the results for the two lower St speeds, for the 0.4 ts/sec St (right panel), both the update-rate-combination and the update-rate-combination \times log (Co/St) effects were highly significant, $F_{3,21} = 32.32, p < .0001$, G-G $p < .0005$ and $F_{18,126} = 6.02, p < .0001$, G-G $p < .001$, respectively. When the update rate of the Co matched that of the St (i.e., when the Co & St update rates were both 60 Hz or both 30 Hz), the psychometric functions were similar (and similar to those for the lower St speeds). However, when the Co and St update-rates differed, the functions were shifted--to the left for the 30 & 60 combination (i.e., when the update rate of the Co was 30 Hz and the update rate of the St was 60 Hz) and to the right for the 60 & 30 combination. Both shifts were statistically significant in individual comparisons of the overall means for the 60 & 60 and 30 & 60 combinations and for the 30 & 30 and 60 & 30 combinations: $F_{1,7} = 16.25, p < .005$, and $F_{1,7} = 24.44, p < .002$, respectively. Because the update-rate-combination \times log (Co/St) interaction was clearly determined in large part by floor and ceiling effects, it was not examined in detail. The log (Co/St) effect was significant at the .0001 level for each of the update-rate combinations.

As shown in Figure 12, despite large individual differences in the magnitude of the update-rate effect for an St of 0.4 ts/sec, the pattern of results is fairly consistent across observers. Although the "unmatched" update-rate functions are not, in all cases, *both* clearly shifted with respect to the "matched" update-rate functions, the 30 & 60 function and the 60 & 30 function are well separated for every observer.




Figure 11
Results of Experiment 1




Figure 12
Update-Rate Effect for Individual Observers

To assess the PSEs, probit analysis was used to fit cumulative normal functions to the 30 & 60 and 60 & 30 data for the 0.4 ts/sec St speed. Because individual response probabilities were based on only 16 trials, two approaches were taken. In one, an analysis was conducted for each of the two update-rate combinations, with each observer treated as a group. In the other, an analysis was conducted for each observer, with each of the four possible update-rate combinations treated as a group. With the former approach, the chi-square goodness-of-fit test and the chi-square parallelism test were both significant in each of the two analyses. For the analysis by observer, both tests were significant for one observer and the parallelism test was significant for another.

Despite the poor fit of the model in some of these analyses, the average estimated PSEs were approximately equal for the two approaches: -0.06 log units for the 30 & 60 condition and 0.09 log units for the 60 & 30 condition. Thus, for the 60-Hz St, a 30-Hz Co that was 13% slower appeared to be of equal speed, whereas for the 30-Hz St, a 60-Hz Co that was 22% faster appeared to be of equal speed. This asymmetry in the magnitude of the shift in PSE was significant ($p < .01$). In addition, the estimated slope for the 30 & 60 combination was higher than that for the 60 & 30 combination. If the magnitude of the perceptual effect of erroneous velocity information increased monotonically with the number of sinusoidal components with the wrong drift velocity, then for the 0.4 ts/sec St, the 7 levels of $\log(\text{Co/St})$ would not have been evenly spaced along the perceptual-speed continuum. The differences in the magnitudes of the shifts and the slopes of the psychometric functions for the 30 & 60 and 60 & 30 combinations are consistent with this possibility.

Given the size of the update-rate effect for the 0.4 ts/sec St, the gradual increase in the prevalence of erroneous sinusoidal components, and the fact that the $\log(\text{Co/St})$ of 0.15 for the 0.2 ts/sec St and the $\log(\text{Co/St})$ of -0.15 for the 0.4 ts/sec St represented the same speed (0.283 ts/sec), it is surprising that an effect of update rate was not apparent for the 0.2 ts/sec St. However, as noted above, the main effect of update rate combination was of borderline significance for this St speed, with the proportion of Co-Faster judgments highest for the 30 & 60 combination and lowest for the 60 & 30 combination. In addition, we compared the average proportion of Co-Faster judgments for the -0.15, -0.10, and -0.05 motion sequences to the average for the 0.05, 0.10, and 0.15 motion sequences for the 60 & 60 and 30 & 30 combinations, first for the 0.2 ts/sec St and then for the 0.4 ts/sec. The difference was significantly larger for the 30 & 30 combination, $F_{1,7} = 6.03$, $p < .05$, suggesting that the increasing amount of erroneous velocity information facilitated velocity discrimination (middle panel, Fig. 11). For the 0.4 ts/sec St, the difference was larger for the 60 & 60 combination, $F_{1,7} = 5.91$, $p < .05$. As shown in the right panel of Figure 12, this effect reflected a difference for the highest $\log(\text{Co/St})$ values, suggesting that when the amount of erroneous velocity information was very high, speed discrimination was more difficult.

In summary, the results of Experiment 1 suggest that, for a point of observation moving a given velocity over a textured ground surface, the perceived speed of self-motion is increased if a large number of visible, sinusoidal components with the wrong drift direction are added to the

optic array. However, because both texture and altitude were constant in this experiment, the speed with respect to the texture pattern and the *global optical flow rate* (Warren, 1982; 1990) covaried. It is therefore possible the effect of update rate on perceived speed depended upon a high flow rate rather than upon a high *texture rate*⁵ (i.e., texture squares per second) and its effect on the temporal frequencies in the original image.

The global optical flow rate is a scalar for (the lower optical hemisphere of) the optic array, indexing the overall speed of the array during constant-velocity, constant-altitude motion over a flat ground (Warren, 1982; 1990). It is defined as the ratio of ground speed to altitude and is typically measured in eye heights (h) per second. In Experiment 1, the global optical flow rates were 12, 24, and 48 h/sec. Although the highest rate could be exceeded during high-speed driving, it is appreciably higher than the flow rates encountered during high-speed, low-level flight. If the effect of update rate is limited to flow rates rarely encountered in simulators, it is of little practical importance.

In order to test the contribution of optical flow rate to the update-rate-combination x St-speed interaction in the first experiment, *optical texture density* (cf., Warren, 1982; 1990) was varied in the second experiment. This was accomplished by modifying the ratio of the altitude of the point of observation to the size of a terrain square. By decoupling these variables, we could present an St with a high texture rate and a relatively low optical flow rate as well as one with a high optical flow rate and a low texture rate. If the update-rate effect in the first experiment was due to a prevalence of visible sinusoidal components drifting in the wrong direction, as has been suggested, then with this texture pattern, update rate should affect perceived speed when the point of observation moves at approximately 0.4 ts/sec, even if the flow rate is low, and update rate should not affect perceived speed when the point of observation moves at less than about 0.2 ts/sec, even if the flow rate is high.

EXPERIMENT 2

Method

The method was the same as that in Experiment 1 with the following exceptions. Six military pilots and 6 nonpilots served as observers. One of the pilots had participated in Experiment 1, and another had participated in exploratory investigations prior to Experiment 1. Three optical texture densities were presented: the original density, 7 times the original density, and one-fourth the original density. The St texture and flow rates for these 3 densities were 0.4 ts/sec and 48 h/s, 0.4 ts/sec and 6.86 h/s, and 0.1 ts/sec and 48 h/s, respectively (see Fig. 13). Thus, the high density St (Fig. 13, middle panel) had the same high texture rate as the original density St (Fig. 13, upper panel), but a lower flow rate, and the low density St (Fig. 13, lower panel) had the same

⁵ We use this term more broadly than it has been used previously (e.g., Denton, 1980; Warren, 1982; Larish & Flach, 1990).




Figure 13
Texture Densities and Optic Flows for Experiment 2

high flow rate as the original density St, but a lower texture rate. (The lower flow rate was equivalent to flying at an altitude of 100 ft at about 400 knots.) Only 5 Co speeds were paired with each St. The log relative speed of the Co ranged from -0.10 to 0.10, in steps of 0.05. The 16 practice trials consisted of the 2 easiest discriminations for the original texture density and speed, presented twice at each of the 4 update-rate combinations. Over 3 test sessions, each of the observers provided 18 judgments for each of the 60 motion-sequence pairs. Each session was divided into 3 blocks of 120 trials. Trials were blocked according to texture density. The order of presentation of the three densities was counterbalanced between and within observers. In each of three subblocks within a block, each of the 20 motion sequence pairs (5 Co speeds x 4 update-rate combinations) was presented twice.

Results and Discussion

The 4-way interaction of group x texture density x update-rate-combination x log (Co/St) is plotted in Figure 14. As suggested by this figure and confirmed by an ANOVA, the performances of the pilot and nonpilot groups did not differ significantly ($p > .20$, for all effects involving group). The main effects of the stimulus factors were all significant, $F_{2,26} = 4.5$, G-G $p < .05$, for texture density; $F_{3,30} = 30.98$, G-G $p < .0001$, for update rate combination; and $F_{4,40} = 209.46$, G-G $p < .0001$, for Log (Co/St), as were their two- and three-way interactions: $F_{6,60} = 28.39$ for texture x update-rate combination; $F_{8,80} = 11.65$ for texture x log(Co/St); $F_{12,120} = 14.27$ for update-rate combination x log(Co/St); and $F_{24,240} = 6.42$ for texture x update-rate combination x log(Co/St), with G-G $p < .0001$, in each case.

In a subsequent ANOVA of the results for the original-texture condition (left panel, Fig. 14), which was characterized by high texture and high flow rates, the update-rate-combination, log Co/St, and update-rate-combination x log Co/St effects were all highly significant (G-G $p < .0001$). Individual contrasts revealed that the average proportion of Co-Faster judgments was higher for the 30 & 60 combination than for the 60 & 60 combination, $F_{1,10} = 27.92$, $p < .0005$, and lower for the 60 & 30 combination than for the 30 & 30 combination, $F_{1,10} = 27.92$, $p < .0005$. The magnitude of these differences varied significantly across levels of log (Co/St). Neither the main effect of update-rate combination nor the interaction of update-rate combination and log (Co/St) was significant in a comparison of the 60 & 60 and 30 & 30 combinations.

For the high-density-texture condition (middle panel, Fig. 14), which was characterized by high texture and low flow rates, the ANOVA results were similar to those for the original-texture condition, with the two main effects and their interaction all highly significant (G-G $p < .0005$). The average proportion of Co-Faster judgments was higher for the 30 & 60 combination than for the 60 & 60 combination, $F_{1,10} = 19.55$, $p < .002$, and lower for the 60 & 30 combination than for the 30 & 30 combination, $F_{1,10} = 25.69$, $p < .0005$. However, in contrast to the results for the original-texture condition, there was a significant difference in the log (Co/St) effect for the 60 & 60 and 30 & 30 combinations, $F_{4,40} = 3.98$, $p < .01$. As shown in Figure 14, the psychometric function for the 30 & 30 condition was steeper than the psychometric function for the 60 & 60 condition. This finding is in accord with that for the 0.2 ts/sec St in Experiment 1.




Figure 14
Results of Experiment 2

For the low-density-texture condition, which was characterized by low texture and high flow rates, there was no evidence of an effect of update-rate combination: Neither the main effect of update-rate combination nor the interaction of update-rate combination and log(Co/St) approached significance ($p > .10$).

It is clear from these results that a high flow rate was neither a necessary nor a sufficient condition for an effect of update rate on the perceived speed of self-motion: An update rate effect was found for a low flow rate and was not found for a high flow rate. The results are consistent with the hypothesis that the update-rate effect is the result of high temporal frequencies in the original image and thus a display image characterized by numerous, visible sinusoidal components drifting in the wrong direction.

As a comparison of the left and middle panels of Figure 14 suggests, however, the difference between the 30 & 60 and 60 & 30 update-rate combinations was greater for the original texture density (moving at 0.4 ts/sec and 48 h/sec) than for the high texture density (moving at 0.4 ts/sec and 6.9 h/sec), $F_{1,10} = 17.46$, $p < .002$. Although this difference could have been due to the difference in flow rate, it seems more likely to have resulted from the difference in texture density itself. With the MIPmap reconstruction filter, the resolution of the texture pattern that was mapped to the far terrain was lower than the resolution of the texture pattern that was mapped to the near terrain. Consequently, for both texture densities, only a portion of the original image contained high temporal frequencies, and thus only a portion of the display image contained numerous sinusoidal components with erroneous drift velocities. This region was *smaller* (and, correspondingly, the region with a predominance of correct velocity information was *larger*) and farther from the central visual field for the high-density texture. In addition, the spatial frequencies of the erroneous components for the high-density texture were higher than those for the original-density texture. The perceptual effect of a sinusoidal component with the wrong drift direction may be a function of its spatial frequency.

This experiment provided no evidence that pilot training or simulator experience affected speed discrimination. Moreover, although individual differences were large, all but one observer showed an effect of update rate and, in most cases, a difference in the size of the effect for the two high-texture rates. Figure 15 presents the data for the original and high texture density conditions for two representative pilots, one who showed very little effect of update rate for either density and one who showed relatively large effects for both.




Figure 15
Response Curves for Two Representative Observers

GENERAL DISCUSSION

The results of this research indicate the update rate of an image generation system may affect the perceived speed of self-motion. The effect has been attributed to the visibility of temporal-sampling-induced spatiotemporal frequencies that provide erroneous velocity information.

In the experiments reported here, the point of observation moved at a constant velocity and altitude over a flat terrain textured with a complex pattern of the type used in modern flight simulators. The spatial spectrum of the texture contained many different frequencies and orientations. Thus, regardless of the speed of the point of observation or the update rate of the IG, the spatiotemporal-frequency spectrum of the display image was characterized by many components with the *correct* velocity information. All of the spatiotemporal frequencies introduced by sampling provided erroneous velocity information.

If the window of visibility, or bandpass, of the human visual system extended to about 30 Hz for this display, then with a 30-Hz update rate, the images were always characterized by visible, temporal-sampling-induced spatiotemporal frequencies, whereas with a 60-Hz update rate, the images for the two lower speeds were free of such frequencies. However, a clear effect of update rate was limited to the highest speed St, for which the spectral overlap was more or less complete for the 30-Hz update rate. The 60-Hz refresh rate (for the 30-Hz update rate) served to attenuate spectral components with a temporal frequency near 30 Hz, and the visual system is relatively insensitive to high temporal frequencies, so this finding is not surprising for the slowest St speed. It is surprising, however, that there was not more evidence of an effect of update rate for the 0.2 ts/sec St. This suggests that the maximum temporal frequency of the window of visibility may have been much less than 30 Hz or that the sensitivity to these spatiotemporal frequencies was low enough to render them perceptually insignificant.

Most of the sampling-induced sinusoidal components with a temporal frequency less than 30 Hz drifted in the wrong direction, not just at the wrong speed. Indeed, if the field of view was limited to the near-terrain regions, it was very difficult to judge the direction of motion in the 30-Hz images for the highest simulated speeds. Nonetheless, when the full image was viewed, these components served to increase the apparent speed of the motion sequence. Moreover, the presence of erroneous components had little, if any, effect on velocity discrimination as long as the images being compared were similarly distorted. There is clearly much yet to be learned about velocity computations for time-varying images in which some but not all spatiotemporal-frequency components provide erroneous information and for a perspective projection in which some but not all spatially local areas of the image have erroneous frequencies.

In both of the experiments reported here, the magnitude of the update rate effect for the 0.4 ts/sec conditions varied substantially across observers. Some showed virtually no effect, whereas the response curves of others were almost flat (near 100% for the 30 & 60 combination and near 0% for the 60 & 30 combination). The data of other observers were in good agreement with the average for the group.

These individual differences could have reflected differences in the size of the area to which an observer allocated visual attention. If this area was small, the observer's judgments may have been based solely on the image motion of the far terrain, where the spatial-frequency band of the texture pattern, and therefore the temporal-frequency band of the original image, had been reduced.

Pursuit eye movements may have played a role as well. Although observers were instructed to fixate a central point marked by a cross hair, all may not have done so. Prior research indicates that the pursuit system is unaffected by high-frequency fluctuations in target velocity, such as those that characterize motion sequences with $k > 0$ (Lindholm & Wetzel, 1992; Lindholm, Wetzel, & Askins, 1996). Therefore, if pursuit eye movements occurred, they were probably unaffected by the update rate of the motion sequence. During smooth pursuit of a small moving target, the motion percept is in accord with the extraretinal rather than the retinal motion signal (Lindholm,

1991, 1994). To the extent that this holds during simulated self-motion over a textured terrain, eye movements would have reduced the effect of update rate in the experiments reported here.

With current simulator technology, there is a tradeoff between image-update rate and scene content. It is often considered desirable to reduce the update rate in order to maintain a relatively high level of detail. Our results indicate that for some combinations of ground speed and ground texture, reducing the update rate to 30 Hz will result in an increase in the perceived speed of self motion. This effect can be eliminated by lowering the rate at which a given texture pattern is traversed (i.e., by reducing ground speed or by mapping the texture pattern to a larger ground area) or by limiting the spatial frequencies in the texture pattern. Both of these solutions restrict the range of temporal frequencies in the original image.

The results also have implications for current efforts to increase near-terrain detail by the application of a "microtexture." These texture patterns add high spatial frequencies to the region of the image with the highest flow rates, potentially resulting in temporal frequencies much higher than half the update rate, at even moderate velocities. Because the projected spatial frequencies of a microtexture will typically be much higher than those of the primary texture, the perceptual consequences of an inadequate update rate may not correspond to the one reported here. However, if erroneous velocity information is provided by many of the components representing a microtexture, self-motion perception is likely to be affected.

REFERENCES

Bracewell, R. N. (1986). *The Fourier transform and its applications*. New York: McGraw-Hill.

Denton, G. G. (1980). The influence of visual pattern on perceived speed. *Perception*, **9**, 393-402.

Finney, D. J. (1971). *Probit analysis*. Cambridge: Cambridge University Press.

Foley, J. D., van Dam, A., Feiner, S. K., & Hughes, J. F. (1992). *Computer graphics: Principles and practice*. Reading, PA: Addison-Wesley.

Gibson, J. J. (1950). *The perception of the visual world*. Boston: Houghton Mifflin.

Gibson, J. J. (1966). *The senses considered as perceptual systems*. Boston: Houghton Mifflin.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.

Gonzalez, R. C., & P. Wintz (1987). *Digital image processing*, Reading, MA: Addison-Wesley.

Graham, N. V. S. (1989). *Visual pattern analyzers*. New York: Oxford University Press.

Heckbert, P. S. (1986). Survey of texture mapping. *IEEE Computer Graphics and Applications*, **6**, 56-67.

Hempstead, C. F. (1966). Motion perception using oscilloscope display. *IEEE Spectrum*, 128-135.

Hochberg, J. (1986). Representation of motion and space in video and cinematic displays. In K. R. Boff, L. K. Kaufman, J. P. Thomas (Eds.) *Handbook of perception and human performance, Vol. 1. Sensory processes and perception* (Ch. 22). New York: Wiley.

Hsu, S. C. (1985). Motion-induced degradations of temporally sampled images, Unpublished Master's Thesis, Mass. Inst. Tech., Cambridge, MA.

Jain, A. K. (1989). *Fundamentals of Digital Image Processing*. Englewood Cliffs, NJ: Prentice-Hall.

Larish, J. F., & Flach, J. M. (1990). Sources of optical information useful for perception of speed of rectilinear self-motion. *J. of Exper. Psych: HP&P*, **16**, 295-302.

Lindholm, J. M. (1991). Perceptual consequences of the filtering characteristics of the pursuit system. *Bulletin of the Psychonomic Society*, 29, 477.

Lindholm, J. M. (1992a). Perceptual effects of spatiotemporal sampling. In M. A. Karim (Ed.) *Electro-optical displays* (pp. 787-807). New York: Marcel Dekker.

Lindholm, J. M. (1992b). *Temporal and spatial factors affecting the perception of computer-generated imagery* (AL-TR-1991-0140). Williams Air Force Base, AZ: Air Force Human Resources Laboratory.

Lindholm, J. M. (1994). Form and motion perception during smooth pursuit. *Investigative Ophthalmology & Visual Science*, 35, 1276.

Lindholm, J. M., & Wetzel, P. A. (1992). Temporal integration characteristics of the smooth pursuit system. *Investigative Ophthalmology & Visual Science*, 33, 1359.

Lindholm, J. M., Wetzel, P. A., & Askins, T. M. (1996). *The stimulus for smooth pursuit* (AL/HR-TR-1996-xx). Mesa, AZ: Armstrong Laboratory, Aircrew Training Research Division.

Sedgwick, H. A. (1986). Space perception. In K. R. Boff, L. K. Kaufman, J. P. Thomas (Eds.) *Handbook of perception and human performance, Vol. 1. Sensory processes and perception* (Ch. 21). New York: Wiley.

Stanley, W. D., Dougherty, G. R., & Dougherty, R. (1984). *Digital Signal Processing*. Reston, VA: Reston.

Stenger, A. J., Zimmerlin, T. A., Thomas, J. P., & Braunstein, M. (1981). *Advanced computer image generation techniques exploiting perceptual characteristics*, AFHRL-TR-80-61, AD-A103 365. Williams Air Force Base, AZ: Operations Training Division, Air Force Human Resources Laboratory.

Warren, R. (1982). *Optical transformation during movement: Review of the optical concomitants of egomotion* (AFOSR-81-0108). Columbus, OH: Ohio State University, Department of Psychology, Aviation Psychology Laboratory.

Warren, R. (1990). Preliminary questions for the study of egomotion. In R. Warren & A. H. Wertheim (Eds.), *The perception and control of self motion* (pp. 3-32). Hillsdale, NJ: Erlbaum.

Watson, A. B., Ahumada, A. J., Jr., & Farrell, J. E. (1986). Window of visibility: A psychophysical theory of fidelity in time-sampled visual motion displays. *Journal of the Optical Society of America*, A3, 300-307.

Watson, A. B., & Ahumada, A. J., Jr., (1985). Model of human visual-motion sensing. *Journal of the Optical Society of America*, A2, 322-342.

Watt, A. (1993). *3D computer graphics*. Reading, MA: Addison-Wesley.